arXiv:1512.06736v1 [cs.DS] 21 Dec 2015

On Lagrangian Relaxation and Reoptimization Problems *

Ariel Kulik † Hadas Shachnai‡ Gal Tamir

December 22, 2015

Abstract

We prove a general result demonstrating the power of Lagrangian relaxation in solving constrained maximization problems with arbitrary objective functions. This yields a unified approach for solving a wide class of *subset selection* problems with linear constraints. Given a problem in this class and some small $\varepsilon \in (0,1)$, we show that if there exists an r-approximation algorithm for the Lagrangian relaxation of the problem, for some $r \in (0,1)$, then our technique achieves a ratio of $\frac{r}{r+1} - \varepsilon$ to the optimal, and this ratio is tight.

The number of calls to the r-approximation algorithm, used by our algorithms, is *linear* in the input size and in $\log(1/\varepsilon)$ for inputs with cardinality constraint, and polynomial in the input size and in $\log(1/\varepsilon)$ for inputs with arbitrary linear constraint. Using the technique we obtain (re)approximation algorithms for natural (reoptimization) variants of classic subset selection problems, including real-time scheduling, the *maximum generalized assignment problem (GAP)* and maximum weight independent set.

1 Introduction

Lagrangian relaxation is a fundamental technique in combinatorial optimization. It has been used extensively in the design of approximation algorithms for a variety of problems (see e.g.,[12, 11, 18, 16, 17, 5] and a comprehensive survey in [20]). In this paper we prove a general result demonstrating the power of Lagrangian relaxation in solving constrained maximization problems of the following form. Given a universe U, a weight function $w : U \to \mathbb{R}^+$, a function $f : U \to \mathbb{N}$ and an integer $L \geq 1$, we want to solve

$$\Pi: \quad \begin{array}{ll} \max_{s \in U} & f(s) \\ \text{subject to:} & w(s) \leq L. \end{array} \tag{1}$$

We solve Π by finding an efficient solution for the Lagrangian relaxation of Π, given by

$$\Pi(\lambda): \quad \max_{s \in U} f(s) - \lambda \cdot w(s), \tag{2}$$

for some $\lambda \geq 0$.

*A preliminary version of this paper appeared in the Proceedings of the 6th Workshop on Approximation and Online Algorithms, Germany, September 2008. Research supported by the Israel Science Foundation (grant number 1574/10), and by the Ministry of Trade and Industry MAGNET program through the NEGEV Consortium.

†Computer Science Department, Technion, Haifa 32000, Israel. E-mail: {kulik,hadas,galtamir}@cs.technion.ac.il

‡Corresponding author.

A traditional approach for using Lagrangian relaxation in approximation algorithms (see, e.g., [11, 16, 5]) is based on initially finding two solutions, SOL_1, SOL_2, for $\Pi(\lambda_1), \Pi(\lambda_2)$, respectively, for some λ_1, λ_2, such that each of the solutions is an approximation for the corresponding Lagrangian relaxation; while one of these solutions is feasible for Π (i.e., satisfies the weight constraint), the other is not. A main challenge is then to find a way to *combine* SOL_1 and SOL_2 to a feasible solution that yields a good approximation for Π. We prove (in Theorem 2.2) a general result, which allows to obtain a solution for Π based on *one* of the solutions only. In particular, we show that, with appropriate selection of the parameters λ_1, λ_2 in the Lagrangian relaxation, we can obtain solutions SOL_1, SOL_2 such that one of them yields an efficient approximation for our original problem Π. The resulting technique leads to fast and simple approximation algorithms for a wide class of *subset selection* problems with linear constraints.

1.1 Subset Selection Problems

Subset selection problems form a large class encompassing such NP-hard problems as real-time scheduling, the generalized assignment problem (GAP) and maximum weight independent set, among others. In these problems, a subset of elements satisfying certain properties needs to be selected out of a universe, so as to maximize some objective function.[1] We apply our general technique to obtain efficient approximate solutions for the following natural variants of some classic subset selection problems.

Budgeted Real Time Scheduling (BRS): The input is a set $\mathcal{A} = \{A_1, \ldots, A_m\}$ of *activities*, where each activity consists of a set of *instances*; an instance $\mathcal{I} \in A_i$ is defined by a half open time interval $[s(\mathcal{I}), e(\mathcal{I}))$ in which the instance can be scheduled ($s(\mathcal{I})$ is the start time, and $e(\mathcal{I})$ is the end time), a cost $c(\mathcal{I}) \in \mathbb{N}$, and a profit $p(\mathcal{I}) \in \mathbb{N}$. A schedule is *feasible* if it contains at most one instance of each activity, and for any $t \geq 0$, at most one instance is scheduled at time t. The goal is to find a feasible schedule, in which the total cost of all the scheduled instances is bounded by a given budget $L \in \mathbb{N}$, and the total profit of the scheduled instances is maximized. *Budgeted continuous real-time scheduling (BCRS)* is a variant of this problem where each instance is associated with a *time window* $\mathcal{I} = [s(\mathcal{I}), e(\mathcal{I}))$ and length $\ell(\mathcal{I})$. An instance $\mathcal{I}$ can be scheduled at any time interval $[\tau, \tau + \ell(\mathcal{I}))$, such that $s(\mathcal{I}) \leq \tau \leq e(\mathcal{I}) - \ell(\mathcal{I}))$. BRS and BCRS arise in many scenarios in which we need to schedule activities subject to resource constraints, e.g., storage requirements for the outputs of the activities.

Budgeted Generalized Assignment Problem (BGAP): The input is a set of bins (of arbitrary capacities) and a set of items, where each item has a size, a value and a packing cost for each bin. Also, we are given a budget $L \geq 0$. The goal is to pack in the bins a feasible subset of items of maximum value, such that the total packing cost is at most L. BGAP arises in many real-life scenarios (e.g., inventory planning with delivery costs).

Budgeted Maximum Weight Independent Set (BWIS): Given a budget L and a graph $G = (V, E)$, where each vertex $v \in V$ has an associated profit p_v (or, *weight*) and associated cost c_v, choose a subset $V' \subseteq V$ such that V' is an *independent set* (i.e., for any $e = (v, u) \in E$, $v \notin V'$ or $u \notin V'$), the total cost of vertices in V', given by $\sum_{v \in V'} c_v$, is bounded by L, and the

[1] We give a formal definition in Section 3.

total profit of V', $\sum_{v \in V'} p_v$, is maximized. BWIS is a generalization of the classical *maximum independent set (IS)* and *maximum weight independent set (WIS)* problems.

1.2 Combinatorial Reoptimization

Traditional combinatorial optimization problems require finding solutions for a single instance. However, many of the real-life scenarios motivating these problems involve systems that change over time. Thus, throughout the continuous operation of such a system, it is required to compute solutions for new problem instances, derived from previous instances. Moreover, since there is some cost associated with the transition from one solution to another, the solution for the new instance must be *close* to the former solution (under certain distance measure).

Solving the resulting *reoptimization* problem involves two challenges: (*i*) *computing* an optimal (or close to the optimal) solution for a new instance, and (*ii*) efficiently *converting* a given solution to a new one (we give below the formal definitions). Indeed, due to the transition costs, we seek for the modified instance an efficient solution which can be reached at low cost. In that sense, the given initial solution plays a restrictive role, rather than serve as guidance to the algorithm.

For example, consider a cloud provider that runs virtual machines for its customers. The provider has a set of servers (hypervisors); each server has a limited resource available for use by the virtual machines that it hosts. The demand for a particular virtual machine changes over time, with a corresponding change in its resource consumption. This may require to migrate some virtual machines among the servers, so as to keep the total demand on each server bounded by its resource availability. Thus, given an initial assignment of virtual machines to the servers, the provider has to find a new assignment which maximizes the profit from serving customers, while minimizing the total migration cost.

1.2.1 Definitions and Notation

A reoptimization problem $R(\Pi)$ is comprised of three elements: a universe U of all feasible solutions, a profit $p : U \to \mathbb{R}^+$ and a transition cost $\delta : U \to \mathbb{R}^+$.

In solving the reoptimization version of a maximization problem, the objective is to find a solution that maximizes the profit, while minimizing the transition cost. In the original optimization problem, we only need to maximize the profit (e.g. $\max\{p(s)|s \in U\}$). We denote this problem by Π and refer to it as the *base problem*. The transition cost for a particular instance of the reoptimization problem represents the cost of moving from an existing solution to a given new solution. (for example, in the cloud provider scenario, this can be the number of virtual machines that need to migrate to reach a new assignment).

In the following, we show how the above notation can be used to describe a reoptimization version of the *maximum spanning tree (MAX-ST)* problem. Denote this problem by Π. Let $G_0 = (V_0, E_0)$ be a weighted graph, and let $T_0 = (V_0, E_{T_0})$ be a MAX-ST for G_0. Let $G = (V, E)$ be a graph derived from G_0 by adding or removing vertices and/or edges, and by (possibly) changing the weights of edges. Let $T = (V, E_T)$ be a MAX-ST for G. For every edge $e \in E_T$, we are given the cost $\delta(e)$ of adding e to the new solution (an example of transition cost $\delta(e)$ can be: $\delta(e) = 0$ if $e \in E_0$; otherwise, $\delta(e) = 1$). The goal in the reoptimization problem R(MAX-ST) is to find a MAX-ST of G with minimal total transition cost. The formal representation of $R(\Pi)$

is $U = \{\text{all spanning trees of } G\}$, $p(T) = \sum_{e \in T} w(e)$, where w is the weight on the graph edges, and $\delta(T) = \sum_{e \in T} \delta(e)$. A polynomial time algorithm for R(MAX-ST) is given in [23].

It is worth noting that the input for the reoptimization problem contains only the new instance of the problem, but not the existing state (T_0 in the MAX-ST example), as the current state is reflected by the transition cost δ.

Approximate Reoptimization: When the problem Π is NP-hard, or when the reoptimization problem $R(\Pi)$ is NP-hard,[2] we seek approximate solutions. The goal is to find a good solution for the new instance, while keeping a low transition cost from the initial solution (or, *configuration*) to the new one. Formally, denote by $\mathcal{O}$ an optimal solution for Π, that is, $p(\mathcal{O}) = \max_{s \in U}(p(s))$. Denote by OPT a solution for $R(\Pi)$ having the minimal transition cost among all the solutions that have a profit $p(\mathcal{O})$. Formally, $\delta(OPT) = \text{argmin}_{s \in U | p(s) = p(\mathcal{O})}(\delta(s))$. We now define the notion of *reapproximation* algorithm.[3]

Definition 1.1 *For $r_1 \geq 1$ and $r_2 \in (0, 1]$ a solution s is an (r_1, r_2)-reapproximation for $R(\Pi)$ if it satisfies (i) $\delta(s) \leq r_1 \cdot \delta(OPT)$, and (ii) $p(s) \geq r_2 \cdot p(OPT)$.*

In this paper we develop a framework that enables to obtain $(1, \alpha)$-reapproximation algorithms for a wide class of subset selection problems, where $\alpha \in (0, 1)$. While some of the resulting approximation ratios may be improved, by applying problem-specific approximation techniques (see Section 1.4), our framework is of interest due to its simplicity and generality. We demonstrate the usefulness of our framework in solving the following reoptimization problem.

The Surgery Room Allocation Problem (SRAP): In a hospital, a surgery room is a vital resource. Operations are scheduled by the severity of patient illness; however, operation schedules tend to change due to sudden changes in patients' condition, the arrival of new patients requiring urgent treatment, or the unexpected absence of senior staff members. Schedule changes involve some costs, e.g., due to the need to rearrange the equipment, or to change the staff members taking care of the patients, as well as their individual schedules.

There is also a profit accrued from each operation. Indeed, some operations are more profitable than others, e.g., due to the coverage received from insurance companies, or due to higher charges in case the operation is scheduled after work hours.

Formally, suppose that the initial input, I_0, consists of n_0 patients. Each patient j is associated with a set $A_{0,j}$ of possible time intervals in which j can be scheduled for operation. An interval $\mathcal{I} \in A_{0,j}$ is a half open time interval $[s_0(\mathcal{I}), e_0(\mathcal{I})$, where $s_0(\mathcal{I}) \leq e_0(\mathcal{I})$. Each interval $\mathcal{I} \in A_{0,j}$ is associated with a profit $p_0(\mathcal{I})$, for all $1 \leq j \leq n_0$. Let S_0 be a given operation schedule for I_0. Consider the input I derived from I_0 by adding or removing patients, by changing the possible time intervals for the patients, or the profits associated with the operations. Suppose that I consists of n patients; each patient j has a set of possible time intervals A_j. Each interval $\mathcal{I} \in A_j$ has a profit $p(\mathcal{I}) \geq 0$ and a transition cost $\delta(\mathcal{I}) \in \mathbb{N}$. This is either the cost of adding $\mathcal{I}$ to the schedule, if $\mathcal{I} \notin S_0$, or the cost of omitting $\mathcal{I}$ from S_0. In any feasible schedule S for I, at most one interval $\mathcal{I} \in A_j$ is selected, for all $1 \leq j \leq n$, and the surgery room is occupied by at most one patient at any time $t \geq 0$. The goal is to find a feasible schedule that maximizes the total profit, while minimizing the aggregate transition cost. In particular, we want to obtain a $(1, \alpha)$-reapproximation algorithm for the problem, for some $\alpha \in (0, 1]$.

[2] As shown in [23], it may be that none, both, or only $R(\Pi)$ is NP-hard.

[3] We refer the reader to [23] for further details.

1.3 Main Results

We prove (in Theorem 2.2) a general result demonstrating the power of Lagrangian relaxation in solving constrained maximization problems with arbitrary objective functions.

We use this result to develop a unified approach for solving subset selection problems with linear constraints. In particular, given a problem Π in this class, and a fixed $\varepsilon \in (0,1)$, we show that if there exists an r-approximation algorithm for the Lagrangian relaxation of Π, for some $r \in (0,1)$, then our technique yields a ratio of $(\frac{r}{r+1} - \varepsilon)$ to the optimal. We further show (in Section 3.4) that this bound is essentially tight, within additive of ε. Specifically, there is a subset selection problem Γ such that, if there exists an r-approximation algorithm for the Lagrangian relaxation of Γ for some $r \in (0,1)$, there is an input I for which finding the solutions SOL_1 and SOL_2 (for the Lagrangian relaxation) and combining these solutions yields at most a ratio of $\frac{r}{r+1}$ to the optimal.

The number of calls to the r-approximation algorithm, used by our algorithms, is *linear* in the input size and in $\log(1/\varepsilon)$, for inputs with cardinality constraint (i.e., where $w(s) = 1$ for all $s \in U$), and polynomial in the input size and in $\log(1/\varepsilon)$ for inputs with arbitrary linear constraint (i.e., arbitrary weights $w(s) \geq 0$).

We apply the technique to obtain efficient approximations for natural variants of some classic subset selection problems. In particular, for the budgeted variants of the real-time scheduling problem we obtain (in Section 4.1) a bound of $(1/3 - \varepsilon)$ for BRS and $(1/4 - \varepsilon)$ for BCRS. For budgeted GAP we give (in Section 4.2) an approximation ratio of $\frac{1-e^{-1}}{2-e^{-1}} - \varepsilon$.

For BWIS we show (in Section 4.3) how an approximation algorithm $\mathcal{A}$ for WIS can be used to obtain an approximation algorithm for BWIS with the same asymptotic approximation ratio. More specifically, let $\mathcal{A}$ be a polynomial time algorithm that finds in a graph G an independent set whose profit is at least $f(n)$ of the optimal, where (i) $f(n) = o(1)$ and (ii) $\log(f(n))$ is polynomial in the input size.[4] Our technique yields an approximation algorithm which runs in polynomial time and achieves an approximation ratio of $g(n) = \Theta(f(n))$. Moreover, $\lim_{n\to\infty} \frac{g(n)}{f(n)} = 1$. Since BWIS generalizes WIS, this implies that the two problems are essentially equivalent in terms of hardness of approximation.

Our technique can be applied iteratively to obtain an $(\frac{r}{1+dr} - \varepsilon)$-approximation algorithm for subset selection problems with $d > 1$ linear constraints, when there exists an r-approximation algorithm for the non-constrained version of the problem, for some $r \in (0,1)$.

It is important to note that the above results, which apply for maximization problems with *linear* constraints, do not exploit the result in Theorem 2.2 in its full generality. We believe that the theorem will find more uses, e.g., in deriving approximation algorithms for subset selection problems with *non-linear* constraints.

Finally, we show how our technique can be used to develop a general $(1,\alpha)$-reapproximation algorithm for any subset selection problem. Specifically, given an instance of a reoptimization problem $R(\Pi)$, where Π is a subset selection problem with an r-approximation algorithm $\mathcal{A}$, we find a $(1, (\frac{r^2}{1+r} - \varepsilon))$-reapproximation algorithm for $R(\Pi)$. We do so by considering a family of *budgeted reoptimization problems* (see Section 5.1) denoted by $R(\Pi, b)$. The problem $R(\Pi, b)$ is a restricted version of $R(\Pi)$, in which we add the constraint that the total transition cost is at most b, for some budget $b \geq 0$. The optimal solution for $R(\Pi, b)$ is denoted $\mathcal{O}_b$. Note that $\mathcal{O}_b$ is the profit of the best solution that can be obtained from the initial solution with

[4] These two requirements hold for most approximation algorithm for the problem.

transition cost at most b. Each of these budgeted sub-problems can be approximated using the Lagrangian relaxation technique. We search for the lowest budget b^* such that $R(\Pi, b^*)$ has a solution of profit that exceeds a certain threshold. Using this algorithm, we derive a $(1, \frac{1}{6} - \varepsilon)$-reapproximation algorithm for SRAP (see Section 5).

1.4 Related Work

Most of the approximation techniques based on Lagrangian relaxation are tailored to handle specific optimization problems. In solving the *k-median* problem through a relation to *facility location*, Jain and Vazirani developed in [16] a general framework for using Lagrangian relaxation to derive approximation algorithms (see also [11]). The framework, that is based on a primal-dual approach, finds initially two approximate solutions SOL_1, SOL_2 for the Lagrangian relaxations $\Pi(\lambda_1)$, $\Pi(\lambda_2)$ of a problem Π, for carefully selected values of λ_1, λ_2; a *convex combination* of these solutions yields a (fractional) solution which uses the budget L. This solution is then rounded to obtain an integral solution that is a good approximation for the original problem. Our approximation technique (in Section 2) differs from the technique of [16] in two ways. First, it does not require rounding a fractional solution: in fact, we do not attempt to combine the solutions SOL_1, SOL_2, but rather, examine each separately and compare the two feasible solutions which can be easily derived from SOL_1, SOL_2, using an efficient transformation of the non-feasible solution, SOL_2, to a feasible one. Secondly, the framework of [16] crucially depends on a primal-dual interpretation of the approximation algorithm for the relaxed problem, which is not required here.

Könemann et al. considered in [17] a technique for solving general partial cover problems. The technique builds on the framework of [16], namely, an instance of a problem in this class is solved by initially finding the two solutions SOL_1, SOL_2 and generating a solution SOL, which combines these two solutions. A comprehensive survey of other work is given in [20].[5]

There has been some earlier work on using Lagrangian relaxation to solve subset selection problems. The paper [21] considered a subclass of the class of the subset selection problems that we study here.[6] Using the framework of [16], the paper claims to obtain an approximation ratio of $r - \varepsilon$ for any problem in this subclass, given a r-approximation algorithm for the Lagrangian relaxation of the problem (satisfying certain properties). Unfortunately, this approximation ratio was shown to be incorrect [22]. Berget et al. considered in [5] the budgeted matching problem and the budgeted matroid intersection problem. The paper gives the first polynomial time approximation schemes for these problems. The schemes, which are based on Lagrangian relaxation, merge the two obtained solutions using some strong combinatorial properties of the problems.

The non-constrained variants of the subset selection problems that we study here are well studied. For known results on real-time scheduling and related problems see, e.g., [2, 7, 3, 4]. Surveys of known results for the generalized assignment problem are given, e.g., in [6, 8, 9, 10].

Numerous approximation algorithms have been proposed and analyzed for the maximum (weight) independent set problem. Alon at al. [1] showed that IS cannot be approximated within factor $n^{-\varepsilon}$ in polynomial time, where $n = |V|$ and $\varepsilon > 0$ is some constant, unless

[5] For conditions under which Lagrangian relaxation can be used to solve discrete/continuous optimization problems see, e.g., [24].

[6] This subclass includes the *real-time scheduling* problem.

$P = NP$. The best known approximation ratio of $\Omega(\frac{\log^2 n}{n})$ for WIS on general graphs is due to Halldórsson [14]. A survey of other known results for IS and WIS can be found e.g., in [13, 15].

To the best of our knowledge, approximation algorithms for the budgeted variants of the above problems are given here for the first time. Our bound for BGAP (in Theorem 4.2) was improved in [19] to $1 - 1/e$.

There is a wide literature on scenarios leading to reoptimization problems, however, most of the earlier studies refer to a model in which the goal is to find an optimal solution for a modified problem instance, with no transition costs incurred in the process (see [23] and the references therein). In this paper, we adopt the reoptimization model introduced in [23]. In this model, it is shown in [23] that for any subset selection problem Π that is polynomially solvable, there is a polynomial time $(1, 1)$-reoptimization algorithm for $R(\Pi)$.

2 Lagrangian Relaxation Technique

Given a universe U, let $f : U \to \mathbb{N}$ be some objective function, and let $w : U \to \mathbb{R}^+$ be a non-negative weight function. Consider the problem Π of maximizing f subject to a budget constraint L for w, as given in (1), and the Lagrangian relaxation of Π, as given in (2).

We assume that the value of an optimal solution s^* for Π satisfies $f(s^*) \geq 1$. For some $\varepsilon' > 0$, suppose that

$$\lambda_2 \leq \lambda_1 \leq \lambda_2 + \varepsilon'. \tag{3}$$

The heart of our approximation technique is the next result.

Theorem 2.1 *For any $\varepsilon > 0$ and λ_1, λ_2 that satisfy (3) with $\varepsilon' = \varepsilon/L$, let $s_1 = SOL_1$ and $s_2 = SOL_2$ be r-approximate solutions for $\Pi(\lambda_1), \Pi(\lambda_2)$, such that $w(s_1) \leq L \leq w(s_2)$. Then for any $\alpha \in [1 - r, 1]$, at least one of the following holds:*

1. $f(s_1) \geq \alpha r f(s^*)$
2. $f(s_2)(1 - \alpha - \varepsilon) f(s^*) \frac{w(s_2)}{L}$.

Proof: Let $L_i = w(s_i)$, $i = 1, 2$, and $L^* = w(s^*)$. From (2) we have that

$$f(s_i) - r f(s^*) \geq \lambda_i (L_i - r L^*). \tag{4}$$

Assume that, for some $\alpha \in [1 - r, 1]$, it holds that $f(s_1) < \alpha r f(s^*)$, then

$$(\alpha - 1) r f(s^*) > f(s_1) - r f(s^*) \geq \lambda_1 (L_1 - r L^*) \geq -r \lambda_1 L^* \geq -r \lambda_1 L.$$

The second inequality follows from (4), the third inequality from the fact that $\lambda_1 L_1 \geq 0$, and the last inequality holds due to the fact that $L^* \leq L$. Using (3), we have

$$\frac{(1 - \alpha) f(s^*)}{L} < \lambda_1 < \lambda_2 + \varepsilon'. \tag{5}$$

Since $\varepsilon' = \varepsilon/L$, we get that

$$\begin{aligned} f(s_2) &\geq \lambda_2 (L_2 - L^*) + r f(s^*) > \left(\frac{(1 - \alpha) f(s^*)}{L} - \varepsilon' \right) (L_2 - L) + r f(s^*) \\ &\geq (1 - \alpha) f(s^*) \frac{L_2}{L} - \varepsilon' L_2 \geq (1 - \alpha - \varepsilon' L) \frac{L_2}{L} f(s^*) = (1 - \alpha - \varepsilon) \frac{L_2}{L} f(s^*) \end{aligned}$$

The first inequality follows from (4), by taking $i = 2$, and the second inequality is due to (5) and the fact that $L^* \leq L$. The third inequality holds since $r \geq 1 - \alpha$, and the last inequality follows from the fact that $f(s^*) \geq 1$. □

We summarize the above discussion in the next theorem, which is the heart of our technique:

Theorem 2.2 *Let $\varepsilon' = \frac{\varepsilon}{L}$.If S_1, S_2 are r-approximation for $\Pi(\lambda_1), \Pi(\lambda_2)$, $\lambda_2 \leq \lambda_1 < \lambda_2 + \varepsilon'$, and $w(S_1) \leq L \leq w(S_2)$, then for all $1 - r \leq \alpha \leq 1$, one of the following holds:*

1. $f(S_1) \geq \alpha r f(S^*)$
2. $f(S_2) \geq (1 - \alpha - \varepsilon) f(S^*) \frac{L_2}{L}$,

where S^* is an optimal solution for Π.

Theorem 2.2 asserts that at least one of the solutions s_1, s_2 is *good* in solving our original problem, Π. If s_1 is a good solution then we have an αr-approximation for Π, otherwise we need to find a way to convert s_2 to a solution s' such that $w(s') \leq L$ and $f(s')$ is a good approximation for Π. Such conversions are presented in Section 3 for a class of *subset selection problems with linear constraints*. Next, we show how to find two solutions which satisfy the conditions of Theorem 2.2.

2.1 Finding the Solutions s_1, s_2

Suppose that we have an algorithm $\mathcal{A}$ which finds a r-approximation for $\Pi(\lambda)$, for any $\lambda \geq 0$. Given an input I for Π, denote the solution which $\mathcal{A}$ returns for $\Pi(\lambda)$ by $\mathcal{A}(\lambda)$, and assume that it is sufficient to consider $\Pi(\lambda)$ for $\lambda \in (0, \lambda_{max})$, where $\lambda_{max} = \lambda_{max}(I)$ and $w(\mathcal{A}(\lambda_{max})) \leq L$.

Note that if $w(\mathcal{A}(0)) \leq L$ then $\mathcal{A}(0)$ is a r-approximation for Π; otherwise, there exist $\lambda_1, \lambda_2 \in (0, \lambda_{max})$ such that λ_1, λ_2, and $s_1 = \mathcal{A}(\lambda_1), s_2 = \mathcal{A}(\lambda_2)$ satisfy (3) and the conditions of Theorem 2.2, and λ_1, λ_2 can be easily found using binary search. Each iteration of the binary search requires a single execution of $\mathcal{A}$ and reduces the size of the search range by half. Therefore, after $R = \lceil \log(\lambda_{max}) + \log(L) + \log(\varepsilon^{-1}) \rceil$ iterations, we have two solutions which satisfy the conditions of the theorem.

Note that the values of λ used during the execution of the algorithm can be represented by $O(\log(\lambda_{max}) + \log(L) + \log(\varepsilon^{-1}))$ bits; thus, we keep the problem size polynomial in its original size and in $\log(\varepsilon^{-1})$.

Theorem 2.3 *Given an algorithm $\mathcal{A}$ which outputs an r-approximation for $\Pi(\lambda)$, and λ_{max}, such that $w(\mathcal{A}(\lambda_{max})) \leq L$, an r-approximate solution or two solutions s_1, s_2 which satisfy the conditions of Theorem 2.2 can be found by using binary search. This requires $\lceil \log(\lambda_{max}) + \log(L) + \log(\varepsilon^{-1}) \rceil$ executions of $\mathcal{A}$.*

We note that when $\mathcal{A}$ is a randomized approximation algorithm whose *expected* performance ratio is r, a simple binary search may not output solutions that satisfy the conditions of Theorem 2.2. In this case, we repeat the executions of $\mathcal{A}$ for the same input and select the solution of maximal value. For some pre-selected values $\beta > 0$ and $\delta > 0$, we can guarantee that the probability that any of the used solutions is not a $(r - \beta)$-approximation is bounded by δ. Thus, with appropriate selection of the values of β and δ, we get a result similar to the result in Theorem 2.2. We discuss this case in detail in the full version of the paper.

3 Approximation Algorithms for Subset Selection Problems

In this section we develop an approximation technique for subset selection problems. We start with some definitions and notation. Given a universe U, let $X \subseteq 2^U$ be a domain, and $f : X \to \mathbb{N}$ a set function. For a subset $S \subseteq U$, let $w(S) = \sum_{s \in S} w_s$, where $w_s \geq 0$ is the *weight* of the element $s \in U$.

Definition 3.1 *The problem*

$$\Gamma : \max_{S \in X} f(S)$$
$$\text{subject to: } w(S) \leq L \tag{6}$$

is a subset selection problem with a linear constraint *if X is a lower ideal, namely, if $S \in X$ and $S' \subseteq S$ then $S' \in X$, and f is a linear non-decreasing set function with $f(\emptyset) = 0$.*[7]

Note that subset selection problems with linear constraints are in the form of (1), and the Lagrangian relaxation of any problem Γ in this class is $\Gamma(\lambda) = \max_{S \in X} f(S) - \lambda w(S)$; therefore, the results of Section 2 hold.

Thus, for example, BGAP can be formulated as the following subset selection problem with linear constraint. The universe U consists of all pairs (i, j) of item $1 \leq i \leq n$ and bin $1 \leq j \leq m$. The domain X consists of all the subsets S of U, such that each item appears at most once (i.e., for any item $1 \leq i \leq n$, $|\{(i', j') \in S : i' = i\}| \leq 1$), and the collection of items that appears with a bin j, i.e., $\{i : (i, j) \in S\}$ defines a feasible assignment of items to bin j. It is easy to see that X is indeed a lower ideal. The function f is $f(S) = \sum_{(i,j) \in S} f_{i,j}$, where $f_{i,j}$ is the profit from the assignment of item i to bin j, and $w(S) = \sum_{(i,j) \in S} w_{i,j}$ where $w_{i,j}$ is the size of item i when assigned to bin j.

The Lagrangian relaxation of BGAP is then

$$\max_{S \in X} f(S) - \lambda w(S) = \max_{S \in X} \sum_{(i,j) \in S} (f_{i,j} - \lambda w_{i,j}) \ .$$

The latter can be interpreted as the following instance of GAP: if $f_{i,j} - \lambda w_{i,j} \geq 0$ then set $f_{i,j} - \lambda w_{i,j}$ to be the profit from assigning item i to bin j; otherwise, make item i infeasible for bin j (set the size of item i to be greater than the capacity of bin j).

We now show how the Lagrangian relaxation technique described in Section 2 can be applied to subset selection problems. Given a problem Γ in this class, suppose that $\mathcal{A}$ is a r-approximation algorithm for $\Gamma(\lambda)$, for some $r \in (0, 1)$. To find λ_1, λ_2 and SOL_1, SOL_2, the binary search of Section 2.1 can be applied over the range $[0, p_{max}]$, where

$$p_{max} = \max_{s \in U} f(s) \tag{7}$$

is the maximum profit of any element in the universe U. To obtain the solutions S_1, S_2 which correspond to λ_1, λ_2, the number of calls to $\mathcal{A}$ in the binary search is bounded by $O(\log(\frac{L \cdot p_{max}}{\varepsilon}))$.

[7] For simplicity, we assume throughout the discussion that $f(\cdot)$ is a *linear* function; however, all of the results in this section hold also for the more general case where $f : 2^S \to \mathbb{N}$ is a non-decreasing *submodular* set function, for any $S \in X$.

Given the solutions S_1, S_2 satisfying the conditions of Theorem 2.2, consider the case where, for some $\alpha \in [1-r, 1]$, property 2 (in the theorem) holds. Denote the value of an optimal solution for Γ by $\mathcal{O}$. Given a solution S_2 such that

$$f(S_2) \geq (1-\alpha-\varepsilon)\frac{w(S_2)}{L} \cdot \mathcal{O}, \tag{8}$$

our goal is to find a solution S' such that $w(S') \leq L$ (i.e., S' is valid for Γ), and $f(S')$ is an approximation for $\mathcal{O}$. We show below how S' can be obtained from S_2. We first consider (in Section 3.1) instances with unit weights. We then describe (in Section 3.2) a scheme for general weights. Finally, we give (in Section 3.3) a scheme which yields improved approximation ratio for general instances, by applying enumeration.

3.1 Unit Weights

Consider first the special case where $w_s = 1$ for any $s \in U$ (i.e., $w(S) = |S|$; we refer to (11) in this case as *cardinality constraint*).

Suppose that we have solutions S_1, S_2 which satisfy the conditions of Theorem 2.2, then by taking $\alpha = \frac{1}{1+r}$ we get that either $f(S_1) \geq (\frac{r}{1+r} - \varepsilon)\mathcal{O}$, or $f(S_2) \geq (\frac{r}{1+r} - \varepsilon)\frac{w(S_2)}{L}\mathcal{O}$. If the former holds then we have a $(\frac{r}{1+r} - \varepsilon)$-approximation for the optimum; otherwise, $f(S_2) \geq (\frac{r}{1+r} - \varepsilon)\frac{w(S_2)}{L}\mathcal{O}$. To obtain S', select the L elements in S_2 with the highest profits.[8] It follows from (8) that $f(S') \geq (1-\alpha-\varepsilon) \cdot \mathcal{O} = (\frac{r}{1+r} - \varepsilon)\mathcal{O}$. Combining the above with the result of Theorem 2.3, we get the following.

Theorem 3.1 *Given a subset selection problem Γ with unit weights, an algorithm $\mathcal{A}$ which yields a r-approximation for $\Gamma(\lambda)$ and λ_{max}, such that $w(\mathcal{A}(\lambda_{max})) \leq L$, a $(\frac{r}{r+1} - \varepsilon)$-approximation for Γ can be derived by using $\mathcal{A}$ and selecting among S_1, S' the set with highest profit. The number of calls to $\mathcal{A}$ is $O(\log(\frac{L \cdot p_{max}}{\varepsilon}))$, where p_{max} is given in (7).*

3.2 Arbitrary Weights

For general element weights, we may assume w.l.o.g. that, for any $s \in U$, $w_s \leq L$. We partition S_2 to a collection of up to $\frac{2W(S_2)}{L}$ disjoint sets $T_1, T_2, \ldots$ such that $w(T_i) \leq L$ for all $i \geq 1$. A simple way to obtain such sets is by adding elements of S_2 in arbitrary order to T_i as long as we do not exceed the budget L. A slightly more efficient implementation has a running time that is linear in the size of S_2 (details omitted).

Lemma 3.2 *Suppose that S_2 satisfies (8) for some $\alpha \in [1-r, 1]$, then there exists $i \geq 1$ such that $f(T_i) \geq \frac{1-\alpha-\varepsilon}{2} \cdot \mathcal{O}$.*

Proof: Clearly, $f(T_1) + \ldots + f(T_N) = f(S_2)$, where $N \leq \frac{2w(S_2)}{L}$ is the number of disjoint sets. By the pigeon hole principle there exists $1 \leq i \leq N$ such that $f(T_i) \geq \frac{f(S_2)}{N} \geq \frac{L \cdot f(S_2)}{2w(S_2)} \geq \frac{1-\alpha-\varepsilon}{2} \cdot \mathcal{O}$. □

Assuming we have solutions S_1, S_2 which satisfy the conditions of Theorem 2.2, by taking $\alpha = \frac{1}{1+2r}$ we get that either $f(S_1) \geq (\frac{r}{1+2r} - \varepsilon)\mathcal{O}$, or $f(S_2) \geq (\frac{2r}{1+2r} - \varepsilon)\frac{w(S_2)}{L}\mathcal{O}$ and can be

[8] When f is a submodular function, iteratively select the element $s \in S_2$ which maximizes $f(T \cup \{s\})$, where T is the subset of elements chosen in the previous iterations.

converted to S' (by setting $S' = T_i$ for T_i which maximizes $f(T_i)$), such that $f(S') \geq (\frac{r}{1+2r} - \varepsilon)\mathcal{O}$, i.e., we get a $(\frac{r}{1+2r} - \varepsilon)$-approximation for Γ.

Combining the above with the result of Theorem 2.3, we get the following.

Theorem 3.3 *Given a subset selection problem Γ with a linear constraint, an algorithm $\mathcal{A}$ that yields an r-approximation for $\Gamma(\lambda)$, and λ_{max}, such that $w(\mathcal{A}(\lambda_{max})) \leq L$, an $(\frac{r}{2r+1} - \varepsilon)$-approximation for Γ can be obtained using $O(\log(\frac{L \cdot p_{max}}{\varepsilon}))$ calls to $\mathcal{A}$, where p_{max} is given in (7).*

3.3 Improving the Bounds via Enumeration

In this section we present an algorithm that uses enumeration to obtain a new problem, for which we apply our Lagrangian relaxation technique. This enables to improve the approximation ratio in Section 3.2 to match the bound obtained for unit weight inputs (in Section 3.1).[9]

For some $k \geq 1$, our algorithm initially 'guesses' a subset T of (at most) k elements with the highest profits in some optimal solution. Then, an approximate solution is obtained by adding elements in U, whose values are bounded by $f(T)/|T|$. Given a subset $T \subseteq U$, we define Γ_T, which can be viewed as the sub-problem that 'remains' from Γ once we select T to be the initial solution. Thus, we refer to Γ_T below as the *residual problem with respect to T*. Let

$$X_T = \left\{ S \;\middle|\; S \cap T = \emptyset, S \cup T \in X, \text{ and } \forall s \in S : f(\{s\}) \leq \frac{f(T)}{|T|} \right\} \tag{9}$$

Consider the residual problem Γ_T and its Lagrangian relaxation $\Gamma_T(\lambda)$:

Γ_T	maximize $f(S)$	$\Gamma_T(\lambda)$	maximize $f(S) - \lambda w(S)$
subject to :	$S \in X_T$	subject to :	$S \in X_T$
	$w(S) \leq L - w(T)$		

In all of our examples, the residual problem Γ_T is a smaller instance of the problem Γ, and therefore, its Lagrangian relaxation is an instance of the Lagrangian relaxation of the original problem. Assume that we have an approximation algorithm $\mathcal{A}$ that, given λ and a pre-selected set $T \subseteq U$ of at most k elements, for some constant $k > 1$, returns an r-approximation for $\Gamma_T(\lambda)$ in polynomial time (if there is a feasible solution for Γ_T). Consider the following algorithm, in which we take $k = 2$:

[9] The running time when applying enumeration depends on the size of the universe (which may be super-polynomial in the input size; we elaborate on that in Section 4.1).

Algorithm 1 *General approximation algorithm*

1. *For any $T \subseteq U$ such that $|T| \leq k$, find solutions S_1, S_2 (for $\Gamma_T(\lambda_1), \Gamma_T(\lambda_2)$ respectively) satisfying the conditions of Theorem 2.2 with respect to the problem Γ_T. Evaluate the following solutions:*

   (a) *$T \cup S_1$*

   (b) *Let $S' = \emptyset$, add elements to S' in the following manner:*
   Find an element $x \in S_2 \backslash S'$ which maximizes the ratio $\frac{f(\{x\})}{w_x}$. If $w(S' \cup \{x\}) \leq L - w(T)$ then add x to S' and repeat the process, otherwise return $S' \cup T$ as a solution.

2. *Return the best of the solutions found in Step 1.*

Let $\mathcal{O} = f(S^*)$ be an optimal solution for Γ, where $S^* = \{x_1, \ldots, x_h\}$. Order the elements in S^* such that $f(\{x_1\}) \geq f(\{x_2\}) \geq \ldots \geq f(\{x_h\})$.

Lemma 3.4 *Let $T_i = \{x_1, \ldots, x_i\}$, for some $1 < i \leq h$, then for any $j > i$, $f(\{x_j\}) \leq \frac{f(T_i)}{i}$.*

In analyzing our algorithm, we consider the iteration in which $T = T_k$. Then $S^* \setminus T_k$ is an optimal solution for Γ_T (since $S^* \setminus T_k \in X_{T_k}$ as in (9)); thus, the optimal value for Γ_{T_k} is at least $f(S^* \backslash T_k) = f(S^*) - f(T_k)$.

Lemma 3.5 *Let S' be the set generated from S_2 by the process in Step 1(b) of the algorithm. Then $f(S') \geq f(S_2)\frac{L-w(T)}{w(S_2)} - \frac{f(T)}{|T|}$*

Proof: Note that the process cannot terminate when $S' = S_2$ since $w(S_2) > L - w(T)$. Consider the first element x that maximized the ratio $\frac{f(\{x\})}{w_x}$, but was not added to S', since $w(S' \cup \{x\}) > L - w(T)$. By the linearity of f, it is clear that

(i) $\frac{f(S' \cup \{x\})}{w(S' \cup \{x\})} \geq \frac{f(\{x\})}{w_x}$, and

(ii) For any $y \in S_2 \backslash (S' \cup \{x\})$, $\frac{f(\{y\})}{w_y} \leq \frac{f(\{x\})}{w_x}$.

Thus, we get that for any $y \in S_2 \backslash (S' \cup \{x\})$, $\frac{f(\{y\})}{w_y} \leq \frac{f(S' \cup \{x\})}{w(S' \cup \{x\})}$, and

$$f(S_2) = f(S' \cup \{x\}) + \sum_{y \in S_2 \backslash (S' \cup \{x\})} f(\{y\}) \leq f(S' \cup \{x\}) \frac{w(S_2)}{w(S' \cup \{x\})}.$$

By the linearity of f, we get that $f(S') + f(\{x\}) = f(S' \cup \{x\}) \geq f(S_2)\frac{L-w(T)}{w(S_2)}$. Since $x \in S_2 \in X_T$, we get $f(\{x\}) \leq \frac{f(T)}{|T|}$. Hence $f(S') \geq f(S_2)\frac{L-w(T)}{w(S_2)} - \frac{f(T)}{|T|}$.

□

Consider the iteration of Step 1. in the above algorithm, in which $T = T_2$ (assuming there are at least two elements in the optimal solution; else $T = T_1$), and the values of the solutions found in this iteration. By Theorem 2.2, taking $\alpha = \frac{1}{1+r}$, one of the following holds:

1. $f(S_1) \geq \frac{r}{1+r}[f(S^*) - f(T)]$

2. $f(S_2) \geq (1 - r - \varepsilon)[f(S^*) - f(T)]\frac{w(S_2)}{L - w(T)}$.

If 1. holds then we get $f(S_1 \cup T) \geq f(T) + (\frac{r}{1+r} - \varepsilon)[f(S^*) - f(T)] \geq (\frac{r}{1+r} - \varepsilon)f(S^*)$, else we have that $f(S_2) \geq (\frac{r}{1+r} - \varepsilon)[f(S^*) - f(T)]\frac{w(S_2)}{L-w(T)}$, and by Lemma 3.5,

$$f(S') \geq f(S_2)\frac{L - w(T)}{w(S_2)} - \frac{f(T)}{|T|} \geq (\frac{r}{1+r} - \varepsilon)[f(S^*) - f(T)] - \frac{f(T)}{|T|}.$$

Hence, we have

$$\begin{aligned} f(S' \cup T) &= f(S') + f(T) \geq f(T) + (\frac{r}{1+r} - \varepsilon)[f(S^*) - f(T)] - \frac{f(T)}{|T|} \\ &= (1 - \frac{1}{k})f(T) + (\frac{r}{1+r} - \varepsilon)[f(S^*) - f(T)] \geq (\frac{r}{1+r} - \varepsilon)f(S^*). \end{aligned}$$

The last inequality follows from choosing $k = 2$, and the fact that $\frac{1}{2} \geq \frac{r}{1+r} - \varepsilon$.

Theorem 3.6 *Algorithm 1 outputs an $(\frac{r}{1+r} - \varepsilon)$-approximation for Γ. The number of calls to algorithm $\mathcal{A}$ is $O((\log(p_{max}) + \log(L) + \log(\varepsilon^{-1}))n^2)$, where $n = |U|$ is the size of the universe of elements for the problem Γ.*

We summarize the above discussion in the next result.

Corollary 3.7 *Given a subset selection problem Γ with a linear constraint, an algorithm $\mathcal{A}$ that yields an r-approximation for $\Gamma(\lambda)$, and λ_{max}, such $w(\mathcal{A}(\lambda_{max})) \leq L$, there is an $(\frac{r}{1+r} - \varepsilon)$-approximation algorithm, such that the number of calls of the algorithm to $\mathcal{A}$ is polynomial in ε, the size of the universe, $|U|$, and the input size.*

In the Appendix we show how our technique for solving subset selection problems with a single linear constraint can be extended to solve such problems with multiple linear constraints, by repetitive usage of our technique.

3.4 Lagrangian Relaxation: Example

We now show the tightness of the bound in Theorem 3.6. Consider the following problem. We are given a base set of elements A, where each element $a \in A$ has a profit $p(a) \in \mathbb{N}$. Also, we have three subsets of elements $A_1, A_2, A_3 \subseteq A$, and a bound $k > 1$. We need to select a subset $S \subseteq A$ of size at most k, such that $S \subseteq A_1$, or $S \subseteq A_2$, or $S \subseteq A_3$, and the total profit from elements in S is, i.e., $\sum_{a \in S} p(a)$, is maximized. The problem can be easily interpreted as a subset selection problem, by taking the universe to be $U = A$, the domain X consists of all the subsets S of U, such that $S \subseteq A_1$ or $S \subseteq A_2$, or $S \subseteq A_3$. The weight function is $w(S) = |S|$, with the weight bound $L = k$, and the profit of a subset S is given by $f(S) = \sum_{a \in S} p(a)$.

The Lagrangian relaxation of the problem with parameter λ is $\max_{S \in X} f(S) - \lambda w(S)$. Assume that we have an algorithm $\mathcal{A}$ that returns an r-approximation for the Lagrangian relaxation of the problem.

For any $\frac{1}{2} > \delta > 0$ and an integer $k > \frac{1}{r} + 4$, consider the following input:

- $A_1 = \{a_1, \ldots, a_{k-1}, b\}$, where $p(a_i) = \frac{1}{r}$ for $1 \leq i \leq k-1$, and $p(b) = k-1$.
- $A_2 = \{c\}$ where $p(c) = k + \delta$.
- $A_3 = \{d_1, \ldots, d_\ell\}$, where $\ell = \lceil \frac{(1+r)(k-1)}{\delta r} \rceil$, and $p(d_i) = 1 + \delta$ for $1 \leq i \leq \ell$.
- $U = A = A_1 \cup A_2 \cup A_3$, and the set S to be chosen is of size at most k.

Denote the profit from a subset $S \subseteq U$ by $p(S)$, and the profit in the Lagrangian relaxation with parameter λ by $p_\lambda(S)$. Clearly, the subset $S = A_1$ is an optimal solution for the problem, of profit $p(A_1) = (k-1)\frac{1+r}{r}$. Consider the possible solutions algorithm $\mathcal{A}$ returns for different values of λ:

- For $\lambda < 1$: the profit from any subset of A_1 is bounded by the original profit of A_1, given by $p(A_1) = (k-1)\frac{1+r}{r}$; the profit from the set $S = A_3$ is equal to $p(A_3) = (1 + \delta - \lambda)\ell \geq \delta\ell \geq (k-1)\frac{(1+r)}{r}$, i.e., A_3 has higher profit than A_1.
- For $1 \leq \lambda \leq \frac{1}{r}$: the profit from any subset of A_1 is bounded by the total profit of A_1 (all elements are of non-negative profit in the relaxation). Taking the difference, we have

$$\begin{aligned} r \cdot p_\lambda(A_1) - p_\lambda(A_2) &= r\left(k - 1 - \lambda + (k-1)(\frac{1}{r} - \lambda)\right) - (k - \lambda) \\ &= rk - r - r\lambda + k - -r\lambda k + rk - k + \lambda \\ &= (1-\lambda)(rk-1) - r \leq 0 \end{aligned}$$

  This implies that in case the optimal set is A_1 (or a subset of A_1), the algorithm $\mathcal{A}$ may choose the set A_2.
- For $\lambda > \frac{1}{r}$: the maximal profit from any subset of A_1 is bounded in this case by $\max\{k - 1 - \lambda, 0\}$, whereas the profit from A_2 is $\max\{k - \lambda, 0\}$.

From the above, we get that $\mathcal{A}$ may return a subset of A_2 or A_3 for any value of λ. However, no combination of elements of A_2 and A_3 yields a solution for the original problem of profit greater than $k(1+\delta)$. This means that, by combining the solutions returned by the Lagrangian relaxation, one cannot achieve approximation ratio better than $\frac{k(1+\delta)}{(k-1)(1+\frac{1}{r})} = \frac{r}{1+r} \cdot \frac{k(1+\delta)}{k-1}$. Since $\frac{r}{1+r} \cdot \frac{k(1+\delta)}{k-1} \to \frac{r}{1+r}$ for $(k, \delta) \to (\infty, 0)$, one cannot achieve approximation ratio better than $\frac{r}{1+r}$.

4 Applications to Budgeted Subset Selection

In this section we show how the technique of Section 3 can be applied to obtain approximation algorithms for several classic subset selection problems with a linear constraint.

4.1 Budgeted Real Time Scheduling

The budgeted real-time scheduling problem can be interpreted as the following subset selection problem with linear constraint. The universe U consists of all instances associated with the activities $\{A_1, \ldots, A_m\}$. The domain X is the set of all feasible schedules; for any $S \in X$, $f(S)$ is

the profit from the instances in S, and $w(S)$ is the total cost of the instances in S (note that each instance is associated with specific time interval). The Lagrangian relaxation of this problem is the classic *interval scheduling* problem discussed in [2]: the paper gives a $\frac{1}{2}$-approximation algorithm, whose running time is $O(n \log n)$, where n is the total number of instances in the input. Clearly, $p_{\max}$ (as defined in (7)) can be used as $\lambda_{\max}$. By Theorem 2.3, we can find two solutions S_1, S_2 which satisfy the conditions of Theorem 2.2 in $O(n \log(n) \log(Lp_{\max}/\varepsilon))$ steps. Then, a straightforward implementation of the technique of Section 3.1 yields a $\left(\frac{1}{3} - \varepsilon\right)$-approximation algorithm whose running time is $O(n \log(n) \log(Lp_{\max}/\varepsilon))$ for inputs where all instances have *unit* cost. The same approximation ratio can be obtained in $O(n^3 \cdot \log(n) \log(Lp_{\max}/\varepsilon))$ steps when the instances may have *arbitrary* costs, using Theorem 3.6 (Note that the Lagrangian relaxation of the residual problem with respect to a subset of elements T is also an instance of the interval scheduling problem).

Consider now the continuous case, where each instance within some activity A_i, $1 \leq i \leq m$, is given by a time window. One way to interpret BCRS as a subset selection problem is by setting the universe to be all the pairs of an instance and a time interval in which it can be scheduled. The size of the resulting universe is unbounded: a more careful consideration of all possible start times of any instance yields a universe of exponential size. The Lagrangian relaxation of this problem is known as *single machine scheduling with release times and deadlines*, for which a $(\frac{1}{2} - \varepsilon)$-approximation algorithm is given in [2]. Thus, we can apply our technique for finding two solutions S_1, S_2 for which Theorem 2.2 holds. However, the running time of the algorithm in Theorem 3.6 may be exponential in the input size (since the number of the enumeration steps depends on the size of the universe, which may be exponentially large). Thus, we derive an approximation algorithm using the technique of Section 3.2. We summarize in the next result.

Theorem 4.1 *There is a polynomial time algorithm that yields an approximation ratio of* $(\frac{1}{3} - \varepsilon)$ *for BRS, and the ratio* $\left(\frac{1}{4} - \varepsilon\right)$ *for BCRS.*

Our results also hold for other budgeted variants of problems that appear in [2].

4.2 The Budgeted Generalized Assignment Problem

Consider the interpretation of GBAP as a subset selection problem, as given in Section 3. The Lagrangian relaxation of BGAP (and also of the deduced residual problems) is an instance of GAP, for which the paper [10] gives a $(1 - e^{-1} - \varepsilon)$-approximation algorithm. We can take in Theorem 2.3 $\lambda_{\max} = p_{\max}$, where $p_{\max}$ is defined by (7), and the two solutions S_1, S_2 that satisfy the condition of Theorem 2.2 can be found in polynomial time. Applying the techniques of Sections 3.1 and 3.3, we get the next result.

Theorem 4.2 *There is a polynomial time algorithm that yields an approximation ratio of* $\frac{1-e^{-1}}{2-e^{-1}} - \varepsilon \approx 0.387 - \varepsilon$ *for BGAP.*

A slightly better approximation ratio can be obtained by using an algorithm of [9]. More generally, our result holds also for any constrained variant of the *separable assignment problem (SAP)* that can be solved using a technique of [10].

4.3 Budgeted Maximum Weight Independent Set

BWIS can be interpreted as the following subset selection problem with linear constraint. The universe U is the set of all vertices in the graph, i.e., $U = V$, the domain X consists of all subsets V' of V, such that V' is an independent set in the given graph G. The objective function f is $f(V') = \sum_{v \in V'} p_v$, the weight function is $w(V') = \sum_{v \in V'} c_v$, and the weight bound is L. The Lagrangian relaxation of BWIS is an instance of the classic WIS problem (vertices with negative profits in the relaxation are deleted, along with their edges). Let $|V| = n$, then by Theorem 3.6, given an approximation algorithm $\mathcal{A}$ for WIS with approximation ratio $f(n)$, the technique of Section 3.3 yields an approximation algorithm $\mathcal{A}_I$ for $BWIS$, whose approximation ratio is $\frac{f(n)}{1+f(n)} - \varepsilon$. The running time of $\mathcal{A}_I$ is polynomial in the input size and in $\log(1/\varepsilon)$. If $\log(1/f(n))$ is polynomial, take $\varepsilon = \frac{f(n)}{n}$; the value $\log(1/\varepsilon) = \log(1/f(n)) + \log(n)$ is polynomial in the input size; thus, the algorithm remains polynomial. For this selection of ε, we have the following result.

Theorem 4.3 *Given an $f(n)$-approximation algorithm for WIS, where $f(n) = o(n)$, for any $L \geq 1$ there exists a polynomial time algorithm that outputs a $g(n)$-approximation ratio for any instance of BWIS with the budget L, where $g(n) = \Theta(f(n))$, and $\lim_{n\to\infty} \frac{g(n)}{f(n)} = 1$.*

This means that the approximation ratios of $\mathcal{A}$ and $\mathcal{A}_I$ are asymptotically the same. Thus, for example, using the algorithm of [14], our technique achieves an $\Omega(\frac{\log^2 n}{n})$-approximation for BWIS. Note that the above result holds for any constant number of linear constraints added to an input for WIS, by repeatedly applying our Lagrangian relaxation technique.

5 Reoptimization of Subset Selection Problems

In this section we show how our Lagrangian relaxation technique can be used to obtain $(1, \alpha)$-reapproximation algorithms for subset selection problems, where $\alpha \in (0, 1)$. To this end, we present the notion of *budgeted reoptimization*. Throughout the discussion, we assume that $R(\Pi)$ is the reoptimization version of a maximization problem Π.

5.1 Budgeted Reoptimization

The budgeted reoptimization problem $R(\Pi, b)$ is a restricted version of $R(\Pi)$, in which we add the constraint that the transition cost is at most b, for some budget $b \geq 0$, and the transition function δ. Formally,

$$R(\Pi, b): \quad \begin{array}{ll} \max_{s \in U} & p(s) \\ \text{subject to:} & \delta(s) \leq b. \end{array} \tag{10}$$

The optimal profit for $R(\Pi, b)$ is denoted $p(\mathcal{O}_b)$, where $\mathcal{O}_b$ is the best solution that can be reached from the initial solution with transition cost at most b.

Definition 5.1 *An algorithm $\mathcal{A}$ yields an r-approximation for $R(\Pi, b)$, for $r \in (0, 1]$, if for any reoptimization input I, $\mathcal{A}$ yields a solution s of profit $p(s) \geq r \cdot p(\mathcal{O}_b)$, and transition cost at most b.*

Example: Assume that Π is the 0-1 Knapsack problem. An instance I of Π consists of a bin of capacity B and n items with profits $p_i \geq 0$ and weights $w_i \geq 0$, for $1 \leq i \leq n$. The formal representation of the problem is $U = \{\text{all feasible packings of the knapsack}\}$, and $p(s) = \sum_{i \in s} p_i$. In the reoptimization version of the problem, $R(\Pi)$, each instance I contains also the transition cost of item i, given by $\delta_i \geq 0$, for $1 \leq i \leq n$. In the budgeted reoptimization version, $R(\Pi, b)$, U is restricted to contain solutions having transition cost at most b. Thus, $U = \{s|\ s \text{ is a feasible packing of the bin, and } \delta(s) \leq b\} = \{s|\ w(s) \leq B, \text{ and } \delta(s) \leq b\}$.

For various problems, $R(\Pi, b)$ satisfies the conditions of Corollary 3.7. In particular, given a problem Π, let $\Gamma_b = R(\Pi, b)$, for some $b \geq 0$, and let $\Gamma_b(\lambda)$ be the Lagrangian relaxation of $R(\Pi, b)$, i.e., $\Gamma_b(\lambda) = \max_{s \in U} p(s) - \lambda \cdot \delta(s)$. If $\Gamma_b(\lambda)$ yields an instance of Π then, by Corollary 3.7, an r-approximation algorithm $\mathcal{A}$ for Π, satisfying for certain value of λ: $w(\mathcal{A}) \leq b$, yields an $(\frac{r}{r+1} - \varepsilon)$-approximation for $R(\Pi, b)$. In the following, we show how this can be used to obtain a reaproximation algorithm for $R(\Pi)$.

5.2 Algorithm

An instance of our reoptimization problem $R(\Pi)$ consists of a universe U of n items. Each item i has a non-negative profit p_i, and a transition cost $\delta_i \in \mathbb{N}$. We give below Algorithm 2, which uses approximation algorithms for Π and $R(\Pi, b)$ in solving $R(\Pi)$.

Algorithm 2 *Reapproximating $R(\Pi)$ for an instance I*

1. *For $r_1, r_2 \in (0, 1]$, let $\mathcal{A}$ be an r_1-approximation algorithm for Π, and let $\mathcal{A}_b$ be an r_2-approximation algorithm for $R(\Pi, b)$.*
2. *Approximate $\Pi(I)$ using $\mathcal{A}(I)$:*

$$Z \leftarrow p(\mathcal{A}(I))$$

3. *Use binary search to find a budget $b > 0$ satisfying:*
   (a) $p(\mathcal{A}_b(I)) \geq r_2 \cdot Z$
   (b) $p(\mathcal{A}_{b-1}(I)) < r_2 \cdot Z$
4. *Return $\mathcal{A}_b(I)$*

Theorem 5.1 *Let I be an instance of the reoptimization problem $R(\Pi)$. For $r_1, r_2 \in (0, 1]$, given an r_1-approximation algorithm $\mathcal{A}$ for Π, and an r_2-approximation algorithm $\mathcal{A}_b$ for $R(\Pi, b)$, for all $b \geq 0$, Algorithem 2 yields in polynomial time a $(1, r_1 \cdot r_2)$-reapproximation for $R(\Pi)$.*

Recall that $\mathcal{O}$ is an optimal solution for Π, and OPT is a solution for $R(\Pi)$ having the minimum transition cost, among the solutions of profit $p(\mathcal{O})$. In Section 5.3 we prove the theorem, by showing that the solution, $S_{\mathcal{A}}$, output by Algorithm 2 has the following properties.

(i) The total transition cost of $S_{\mathcal{A}}$ is at most the transition cost of OPT, i.e.,

$$\delta(S_{\mathcal{A}}) \leq \delta(OPT)$$

(ii) The profit of $S_{\mathcal{A}}$ satisfies

$$p(S_{\mathcal{A}}) \geq r_1 \cdot r_2 \cdot p(OPT).$$

Combining Theorem 5.1 and Corollary 3.7, we show that a wide class of reoptimization problems can be approximated using our technique.

Corollary 1 *Let $R(\Pi)$ be the reoptimization version of a subset selection problem Π, and let $\Gamma_b = R(\Pi, b)$, for $b \geq 0$. Denote by $\mathcal{A}$ an r-approximation algorithm for Π, for $r \in (0,1)$. If the lagrangian relaxation of Γ_b, $\Gamma_b(\lambda)$, yields an instance of the base problem Π, for all $b \geq 0$, then for any $\varepsilon > 0$, Algorithm 2 is a $(1, \frac{r^2}{1+r} - \varepsilon)$-reapproximation algorithm for $R(\Pi)$.*

Proof: By Corollary 3.7, given $\varepsilon' > 0$, we have an $(\frac{r}{r+1} - \varepsilon')$-approximation algorithm, $\mathcal{A}_b$, for $R(\Pi, b)$, for any $b \geq 0$. Thus, using Theorem 5.1 with algorithms $\mathcal{A}$ and $\mathcal{A}_b$, and taking $\varepsilon' = \frac{\varepsilon}{r}$, we obtain a $(1, r \cdot (\frac{r}{1+r} - \frac{\varepsilon}{r}))$-reapproximation algorithm for $R(\Pi)$. □

5.3 Proof of Theorem 5.1

We use in the proof the next lemmas.

Lemma 5.2 *The solution output by Algorithm 2 for an instance I satisfies $\delta(S_{\mathcal{A}}) \leq \delta(OPT)$.*

Proof: Let OPT be a solution of minimum transition cost, among those that yield an optimal profit for Π, and let $b^* \geq \delta(OPT)$. By definition, OPT is a valid solution of $\Pi(R, b^*)$. Hence, the optimal profit of $\Pi(R, b^*)$ is $p(OPT)$, and we have

$$p(\mathcal{A}_{b^*}(I)) \geq r_2 \cdot p(OPT) \geq r_2 \cdot Z.$$

It follows that, for any budget b satisfying $\mathcal{A}_b(I) < r_2 \cdot Z$, we have $b < \delta(OPT)$. By Step (3b), the algorithm selects a budget b such that $p(\mathcal{A}_{b-1}(I)) < r_2 \cdot Z$. Hence, $b - 1 < \delta(OPT)$. Since b and $\delta(OPT)$ are integers, we have that $b \leq \delta(OPT)$. □

Lemma 5.3 *The profit of $S_{\mathcal{A}}$ satisfies $p(S_{\mathcal{A}}) \geq r_1 \cdot r_2 \cdot p(OPT)$.*

Proof: In Step (3a), Algorithm 2 selects a solution of profit at least $r_2 \cdot Z$. Also, Z is an r_1-approximation for Π; therefore, $Z \geq r_1 \cdot p(OPT)$. This yields the statement of the lemma. □

Lemma 5.4 *Algorithm 2 has polynomial running time.*

Proof: The algorithm proceeds in three steps. Step 2 is polynomial since $\mathcal{A}$ runs in polynomial time. In Step 3 we search over all budgets $0 \leq b \leq b_{max} = \sum_{a \in I} \delta(a)$. While b_{max} may be arbitrarily large, $log(b_{max})$ is polynomial in the input size, and indeed Algorithm 2 calls $\mathcal{A}_b$ $O(log(b_{max}))$ times. □

Combining the above lemmas, we have the statement of the theorem. □

5.4 A Reapproximation Algorithm for SRAP

We now show how to use Algorithm 2 to obtain a $(1, \alpha)$-reapproximation algorithm for SRAP, for some $\alpha \in (0, 1)$. Recall, that an input for the *real-time scheduling problem* consists of a set $\mathcal{A} = \{A_1, \ldots, A_m\}$ of *activities*, where each activity consists of a set of *instances*; an instance $\mathcal{I} \in A_j$ is defined by a half open time interval $[s(\mathcal{I}), e(\mathcal{I}))$ in which the instance can be scheduled ($s(\mathcal{I})$ is the start time, and $e(\mathcal{I})$ is the end time), and a profit $p(\mathcal{I}) \geq 0$. A schedule is *feasible* if it contains at most one instance of each activity, and for any $t \geq 0$, at most one instance is scheduled at time t. The goal is to find a feasible schedule of a subset of the activities that maximizes the total profit (see, e.g., [2]). Let Π be the real-time scheduling problem. Then SRAP can be cast as $R(\Pi)$, the reoptimization version of Π.

Now, given budgeted SRAP, $\Gamma_b = R(\Pi, b)$, in which the transition cost is bounded by b, for some $b \geq 0$, we can write Γ_b in the form

$$\begin{aligned} \Gamma_b : \ & \max_{S \in X} \ f(S) \\ & \text{subject to: } w(S) \leq b, \end{aligned} \tag{11}$$

where $X = \{\text{all feasible operation schedules}\}$, and $w(S) = \delta(S)$. The Lagrangian relaxation of Γ_b is $\Gamma_b(\lambda) : \ \max_{S \in X} f(S) - \lambda \cdot w(S)$,. We note that $\Gamma_b(\lambda)$ yields an instance of the real-time scheduling problem, Π. Our base problem, Π, can be approximated within factor 1/2 [2]. As shown in Section 4.1, budgeted real-time scheduling admits a $(\frac{1}{3} - \varepsilon)$-approximation. The next result follows from Theorem 5.1.

Theorem 5.5 *There is a polynomial-time* $(1, 1/6 - \varepsilon)$*-reapproximation algorithm for SRAP.*

References

[1] N. Alon, U. Feige, A. Wigderson, and D. Zuckerman. Derandomized graph products. *Computational Complexity*, 5(1):60–75, 1995.

[2] A. Bar-Noy, R. Bar-Yehuda, A. Freund, J. Naor, and B. Schieber. A unified approach to approximating resource allocation and scheduling. *Journal of the ACM*, pages 1–23, 2000.

[3] A. Bar-Noy, S. Guha, J. Naor, and B. Schieber. Approximating the throughput of multiple machines in real-time scheduling. *SIAM J. Comput.*, 31(2):331–352, 2001.

[4] R. Bar-Yehuda, M. Beder, Y. Cohen, and D. Rawitz. Resource allocation in bounded degree trees. *Algorithmica*, 54(1):89–106, 2009.

[5] A. Berger, V. Bonifaci, F. Grandoni, and G. Schäfer. Budgeted matching and budgeted matroid intersection via the gasoline puzzle. In *IPCO*, 2008.

[6] C. Chekuri and S. Khanna. A PTAS for the multiple knapsack problem. *SIAM J. Comput.*, 35(3):713–728, 2006.

[7] J. Chuzhoy, R. Ostrovsky, and Y. Rabani. Approximation algorithms for the job interval selection problem and related scheduling problems. *Math. Oper. Res.*, 31(4):730–738, 2006.

[8] R. Cohen, L. Katzir, and D. Raz. An efficient approximation for the generalized assignment problem. *Inf. Process. Lett.*, 100(4):162–166, 2006.

[9] U. Feige and J. Vondrák. Approximation algorithms for allocation problems: Improving the factor of 1 - 1/e. In *FOCS*, pages 667–676, 2006.

[10] L. Fleischer, M. X. Goemans, V. S. Mirrokni, and M. Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *SODA*, pages 611–620, 2006.

[11] N. Garg. A 3-approximation for the minimum tree spanning k vertices. In *FOCS*, 1996.

[12] M. X. Goemans and R. Ravi. The constrained minimum spanning tree problem. In *SWAT*, pages 66–75, 1996.

[13] M. M. Halldórsson. Approximations of independent sets in graphs. In *APPROX*, 1998.

[14] M. M. Halldórsson. Approximations of weighted independent set and hereditary subset problems. *J. Graph Algorithms Appl.*, 4(1), 2000.

[15] M. M. Halldórsson. Approximations of weighted independent set and hereditary subset problems. *Journal of Graph Algorithms and Applications*, pages 1–16, 2000.

[16] K. Jain and V. V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *Journal of the ACM*, pages 274–296, 2001.

[17] J. Könemann, O. Parekh, and D. Segev. A unified approach to approximating partial covering problems. In *ESA'06: Proceedings of the 14th conference on Annual European Symposium*, pages 468–479, London, UK, 2006. Springer-Verlag.

[18] J. Könemann and R. Ravi. A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees. *SIAM J. Comput.*, 31(6):1783–1793, 2002.

[19] A. Kulik. *Submodular and Linear Maximization with Knapsack Constraints*. MSc Thesis, Computer Science Department, Technion, 2011.

[20] J. Mestre. *Primal-Dual Algorithms for Combinatorial Optimization Problems*. PhD thesis, CS Dept, Univ. of Maryland, 2007.

[21] J. Naor, H. Shachnai, and T. Tamir. Real-time scheduling with a budget. *Algorithmica*, pages 343–364, 2007.

[22] G. Schäfer. *Private communication*, 2007.

[23] H. Shachnai, G. Tamir, and T. Tamir. A theory and algorithms for combinatorial reoptimization. In *LATIN 2012: Theoretical Informatics - 10th Latin American Symposium, Arequipa, Peru, April 16-20, 2012. Proceedings*, pages 618–630, 2012.

[24] L. A. Wolsey and G. L. Nemhauser. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1999.

A Solving Multi-budgeted Subset Selection Problems

In the following we extend our technique as given in Section 3 to handle subset selection problems with d linear constraints, for some $d > 1$. More formally, consider the problem:

$$\begin{aligned} \max_{S \in X} \quad & f(S) \quad \text{subject to:} \\ \forall_{1 \leq i \leq d} : \quad & w_i(S) \leq L_i, \end{aligned} \tag{12}$$

where X is a lower ideal, and the functions f and w_i for $1 \leq i \leq d$ are non-decreasing linear set function, such that $f(\emptyset) = w_i(\emptyset) = 0$. This problem can be interpreted as the following subset selection problem with a single linear constraint. Let $X' = \{S \in X | \forall_{1 \leq i \leq d-1} : \ w_i(S) \leq L_i\}$; the linear constraint is $w_d(S) \leq L_d$, and the function f remains as defined above. The Lagrangian relaxation of (12) has the same form (after removing in the relaxation elements with negative profits), but the number of constraints is now $d - 1$. This implies that, by repeatedly applying the technique in Section 3.3, we can obtain an approximation algorithm for (12) from an approximation algorithm for the non-constrained problem (in which we want to find $\max_{S \in X} f'(S)$, where f' is some linear function). Thus, given an r-approximation algorithm for the problem after "relaxing" d constraints, we derive an $(\frac{r}{1+dr} - \varepsilon)$-approximation algorithm for (12). Note that there is a simple reduction[10] of the problem in (12) to the same problem with $d = 1$, which yields a $\frac{\rho}{d}$-approximation for (12), given a ρ-approximation algorithm $\mathcal{A}$ for the problem with single constraint. For sufficiently small $\varepsilon > 0$, the ratio of $\frac{\rho}{1+(d-1)\rho} - \varepsilon$ obtained by repeatedly applying Lagrangian relaxation and using the approximation algorithm $\mathcal{A}$ is better, for any $\rho \in (0, 1)$.

[10] Assume w.l.o.g that $L_i = 1$ for every $1 \leq i \leq d$, and set the weight of an element e to be $w_e = \max_{1 \leq i \leq d} w_i(\{e\})$